Received: 22 May 2023 | Revised: 13 July 2023 | Accepted: 1 August 2023

DOI: 10.1002/psp4.13027

TUTORIAL

Building an adaptive dose simulation framework to aid dose and schedule selection

Richard Hooijmaijers[1] | **Ridhi Parasrampuria**[2] | **Eleonora Marostica**[1] | **Geraldine Ferron-Brady**[2] | **Teun M. Post**[1] | **Sandra A. G. Visser**[2]

[1]Leiden Experts on Advanced Pharmacokinetics and Pharmacodynamics (LAP&P), Leiden, The Netherlands

[2]GSK, Collegeville, Pennsylvania, USA

Correspondence
Sandra A. G. Visser GSK, Collegeville, PA, USA.
Email: sandra.x.visser@gsk.com

Abstract

Establishing a dosing regimen that maximizes clinical benefit and minimizes adverse effects for novel therapeutics is a key objective for drug developers. Finding an optimal dose and schedule can be particularly challenging for compounds with a narrow therapeutic window such as in oncology. Modeling and simulation tools can be valuable to conduct in silico evaluations of various dosing scenarios with the goal to identify those that could minimize toxicities, avoid unscheduled dose interruptions, or minimize premature discontinuations, which all could limit the potential for therapeutic benefit. In this tutorial, we present a stepwise development of an adaptive dose simulation framework that can be used for dose optimization simulations. The tutorial first describes the general workflow, followed by a technical description with basic to advanced practical examples of its implementation in mrgsolve and is concluded with examples on how to use this in decision-making around dose and schedule optimization. The adaptive simulation framework is built with pharmacokinetic, pharmacodynamic (i.e., biomarkers, activity markers, target engagement markers, efficacy markers), and safety models that include evaluations of unexplained interindividual and intraindividual variability and covariate impact, which can be replaced and expanded (e.g., combination setting, comparator setting) with user-defined models. Subsequent adaptive simulations allow investigation of the impact of starting dose, dosing intervals, and event-driven (exposure or effect) dose modifications on any end point. The resulting simulation-derived insights can be used in quantitatively proposing dose and regimens that better balance benefit and adverse effects for further evaluation, aiding dose selection discussions, and designing dose modification recommendations, among others.

INTRODUCTION

Clinical dose and schedule selection are critically important in the development of novel therapeutics and drug combinations irrespective of the therapeutic area. For some compounds, in particular those that have a wide therapeutic index (i.e., a benefit/risk constant over a reasonable dose/exposure range) a single dose and regimen may be applicable to a large patient population without a need for dose adjustments. However, for other compounds, dose and regimen adjustments may be required for specific patient populations or need to be adjusted during treatment because of adverse events or the need for careful titration of therapeutic effect. These could be compounds with a narrow therapeutic index and also compounds for which the exposure or pharmacology changes significantly as a result of intrinsic or extrinsic factors. Perhaps the most well-known drug with a narrow therapeutic index is warfarin, for which its therapeutic effect (anticoagulation) is monitored regularly based on which patients receive individualized dosing recommendations.[1]

In oncology, targeted therapies with a narrow therapeutic index have often high rates of (unscheduled) dose interruptions and dose reductions to manage adverse events, which could limit the potential for therapeutic benefit.[2] The US Food and Drug Administration's Project Optimus initiative to reform the dose optimization and dose selection paradigm in oncology drug development has brought the challenge of finding a dosing regimen that maximizes clinical benefit and minimizes (chronic) adverse effects to the center of attention.[3–7] The recently released draft guidance on oncology dose optimization recommends characterizing the pharmacokinetics (PK), activity, safety, and tolerability across multiple dosages and that dosages selected for administration in a clinical trial should be adequately supported by data appropriate to the stage of development for each indication and usage.[8]

If emerging clinical data indicate that a particular compound will likely require dose modifications or a withholding of doses during treatment to manage either tolerability or suboptimal efficacy, the use of adaptive dose simulations can be valuable to investigate the benefit of planned schedule changes, intermittent schedules, and adaptive dose titration to balance benefit and chronic adverse effects. This requires the generation of PK, pharmacodynamics (PD), efficacy, and safety models using emerging clinical data across a range of different doses as is typically collected during phase I trials[9] or using a combination of clinical and preclinical data.[10] The simulation insights can be used to address dosing regimen–related questions throughout drug development, whether it is for the design of a new clinical trial or to update dose recommendations in specific patient or treatment groups based on experience in the postmarketing setting (Figure 1).

In this tutorial, we introduce stepwise how an adaptive dose simulation framework can be set up for compounds that could benefit from dose modifications to manage tolerability, address suboptimal efficacy, or target specific exposures. We first introduce the general workflow, followed by a technical description with practical examples of increasing complexity of its implementation in mrgsolve (https://mrgsolve.org/). We conclude by discussing how to integrate the derived knowledge into decisions about dose and dose regimens. It should be noted that this tutorial does not go into statistical criteria, power calculations, or trial enrichment per se,[11] but the framework itself is compatible with these purposes. Neither is individualized precision dosing part of the tutorial, but insights gained using a framework can inform decision-making on whether individualized dosing should be considered.

GENERAL WORKFLOW

Once a need is identified to explore dose and regimens that include event-driven dose modifications during the duration of treatment, a general workflow is proposed that includes (1) engaging a multidisciplinary team to define the objectives that address the key question(s) and align on the framework components, (2) collecting information around the framework components, (3) setting up the technical framework and simulation of scenarios, (4) sharing simulation results addressing the objectives with team. Figure 2 aims to capture this iterative workflow schematically. In the supplemental material (Appendix S1), a checklist of potential questions is provided that can help guide the construction of a specific framework. It should be noted that this framework is generally applicable to any modeling and simulation activity that needs multidisciplinary team collaboration.

Engage the team

The first step in the practical execution of an adaptive dosing framework is the formation of a multidisciplinary team. This team should at a minimum consist of a clinician, clinical pharmacologist, a pharmacometrician, and other relevant members of the project team, such as a statistician to discuss the objectives for the adaptive dose simulations. In this team, the objectives and desired outcome of the adaptive dose simulation framework should be clearly defined. What are the key questions that need to be addressed? Is it to explore dose and schedule options to be included in the design of a new clinical trial, or to

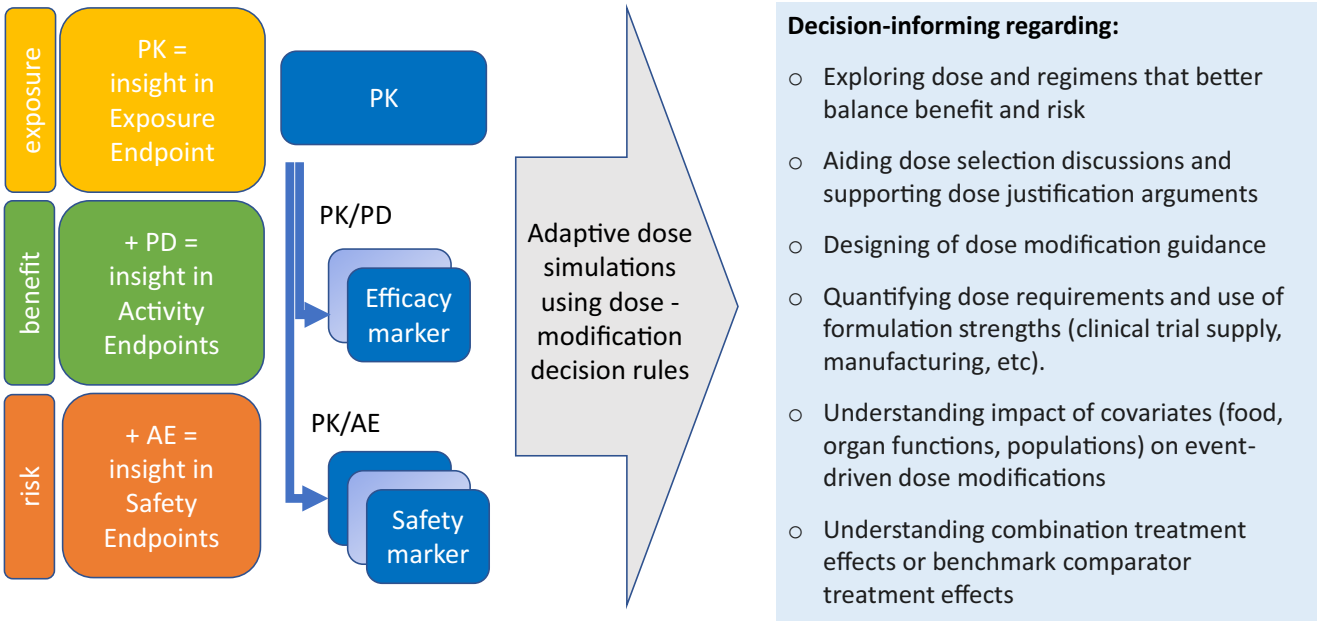

Decision Rules: Exposure-, activity, or safety driven dose adjustments that impact longitudinally the dose, exposure, activity (e.g., efficacy or target engagement) and safety endpoints.

- **Exposure driven:** e.g., Therapeutic Drug Monitoring
- **Activity driven**: e.g., warfarin individualized dosing based on monitoring coagulation
- **Safety driven**: e.g., ruxolitinib dose titration

FIGURE 1 Schematic overview of adaptive dose simulation framework. AE, adverse events; PD, pharmacodynamics; PK, pharmacokinetics.

investigate the impact of starting dose and dose modification recommendation on the overall predicted outcome (efficacy, safety, or both), or to understand the overall compound (manufacturing) demand for a clinical program, or to provide justification for dose and schedule recommendations, or something else? Sometimes the objective discussion may need to be preceded by a more educational session on the benefits and limitations of modeling-and-simulation activities and what valuable insights an adaptive dose simulation framework could deliver.

When the objectives are clear, the next steps are to define the framework components, understand the intended patient population, explore and define decision criteria for dose adaptations, align on clinically feasible dose adaptation options to investigate, and define simulation outcomes that address the questions. This is expected to be an iterative process that may require multiple team meetings; therefore, expectation management needs to be a recurring component of these interactions. It is important to consciously select the characteristics (e.g., demographics, disease state, clinical parameters, covariates, inclusion and exclusion criteria) of the relevant patient group(s) to allow definition of the correct assumptions around the PK, PD (i.e., biomarkers, activity markers, target engagement markers, efficacy markers), and/or safety models to be used in the simulations. In addition, the selection of quantifiable clinical decision criteria and predefined dose adaptation rules are important. Although a large variety of configurations for adaptation criteria and dosing rules can be implemented, subjective criteria at the discretion of the treating clinician cannot. In addition, ad hoc modifications to decision criteria may have a substantial impact on the timing and delivery of results. Clinicians should therefore be encouraged to work out a consensus on the (potential) clinical decision criteria for dose adaptations very thoroughly at an early stage. Note that the selection of different subsets of patients may facilitate resolution of disagreements among clinicians by formalizing situations in which a clinician might be tempted to diverge from a standard treatment protocol for a specific patient. Possible iterations and modifications to decision criteria may be discussed as the amount of knowledge increases based on initial adaptive simulation results and evolving insights from early data. In theory, there is no limit to the number of markers or dosing decisions, but the complexity of implementation and quality control, duration of simulations, and time required to output the results will increase exponentially. Flowcharts are a useful tool to keep track of different decision criteria and dose adaptation pathways. Updates of these charts can be used to keep track of discussions over time, which become more valuable as the objectives increase in complexity. In addition, the team

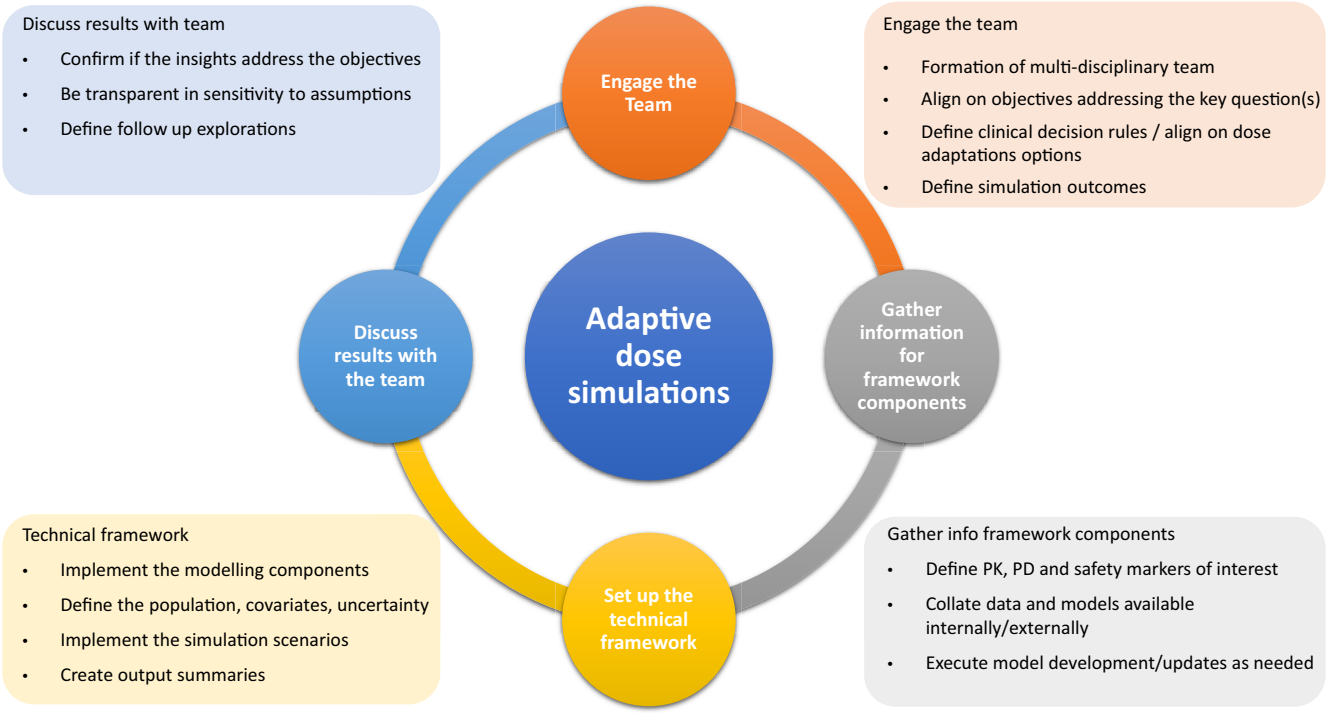

FIGURE 2 General workflow of setting up an adaptive dose simulation framework. This workflow is iterative, as new insights may trigger new questions, additional data generation, new model updates, and new simulations. PD, pharmacodynamics; PK, pharmacokinetics.

should agree on what would be insightful summaries and visualizations of the simulated results. Finally, additional project management-related items should be discussed as well, for example, including whether to prepare an analysis plan, what the expected timelines are, who the project manager will be, and with what frequency additional team meetings should be called. It is advised to start this alignment process early as the combining of all input sources and the discussion process itself can be time-consuming.

Gather information for framework components

The next step involves the preparation of an overview of all components required for the technical implementation that includes defining relevant end points, summarizing available data, models, and options for dose modifications but also potentially analyzing data to develop additional models. The core of this overview is formed by defining the PK, PD (including biomarker, activity, efficacy), and safety data of interest, whose definition and selection can be taken from sources such as study protocols, clinical experience, competitor information and registration dossiers, and conference abstracts and presentations. Based on these, an exploration of available data and models can be done internally and if applicable on publicly available literature. If unavailable or not yet suitable, models may need to be developed or updated. Pharmacometricians are equipped to evaluate the quality of the available models considering the criteria and outcome measures as defined in the team engagements, and the outcome of such an assessment can affect the choice of markers to guide the dosing adaptations. Therefore, in practice an iterative process of team engagement and the technical discussion on data and model quality occurs, gradually refining and updating the plan as more insight into its feasibility and weaknesses is gained. In addition, new information may come in over time from various internal and external sources, which can also prompt reevaluation of criteria, decisions, and/or simulation designs.

Set up the technical framework

Following the team engagement meetings, the technical implementation involves preparation of a list of practical simulation settings (e.g., events, criteria and dosing implementations, including their specific timing, thresholds), with a proper bookkeeping of events. The feedback to the team forms a highly recommended check to see whether the results of earlier discussions were clear and correctly interpreted. In addition, a graphical representation of the available data and/or simulations can be created using available models; such illustrations can be used to assist and guide team

discussions on the limitations of available models and data and on the suitability of assumptions posited in the team meetings. Practical settings for the simulations that result from team discussions also include aspects such as the number of subjects, the number of visits, the duration of simulations, the number of sampling times, the choice to include parameter uncertainty, interindividual variability and/or residual unexplained variability, and other technical aspects. Graphical illustration of the impact of these different choices takes time, but it is recommended to focus team discussions and come to an agreed-upon list of technical specifications. Examples on a step-by-step setup are provided in the Technical Workflow section of this tutorial.

Discuss results with the team

Once the simulations are completed, the results are discussed with the team, in view of their usefulness in addressing the original question(s) for which the framework was designed. In practice, the results may prompt new insights or assumptions that can result in modification of assumptions, simulation parameters, and other aspects, in which case the described steps can be performed once more. The results from basic designs can often be reported using simple tabular summary statistics of, for example, the fraction of patients experiencing a dose adjustment within a particular period, the average dose over the simulated interval, the time to dose adjustment, the summary of benefit (PD, activity, efficacy) and risk (safety) during the treatment duration. However, for more complex designs, a correct and informative presentation of the results may require complex graphical representations, for instance, when wanting to have a comprehensive look at the balance between efficacy and safety of different doses and dose regimens along with dose modifications or when comparing against an existing treatment. Dummy examples summarizing the output of these complex designs can be discussed with the team at an early stage to manage expectations and to start early with the script-based refinement of such plots, which may be time-consuming. Different graphical outputs are compatible with this framework.[12–14]

TECHNICAL WORKFLOW

Basic setup

This tutorial is based on simulations of PK, PD, and safety profiles in R[15] using the open source mrgsolve package.[16] This package allows simulations from hierarchical, closed-form solutions and ordinary differential equation (ODE)–based models that are widely employed in the pharmacometrics community to support drug development. Although alternative implementations of adaptive simulation frameworks are viable (e.g., NONMEM,[17] nlmixr2/rxode2,[18,19] OptiDose[20]), we chose mrgsolve for the following benefits: (1) the C++ – based coding provides fast output compared with alternatives, with well-established and tested integration into R; (2) the free and open source nature facilitates dissemination of this framework and provides opportunities for collaboration and wide applications; (3) the package can be integrated with commonly used R-based packages for data management (`dplyr`), plotting (`ggplot`, `lattice`), and user interface (`Shiny`); and (4) implementation and coding of the model and its parameters are compatible with output from NONMEM.

Using the mrgsolve package requires defining the underlying model, as well as a command section that defines the input conditions and output specifications. The basic command setup and syntax is quite instinctive and comprises the following sections:

```
library(mrgsolve)                                          (1)
mod <- mread_cache("pk1", modlib())                        (2)
mod                                                        (3)
param(mod)                                                 (4)
evnt <- ev(amt = 100, ii = 24, addl = 9)                   (5)
out <- mod %>% ev(evnt) %>% mrgsim(end = 480,
delta = 0.1)                                               (6)
out                                                        (7)
plot(out)                                                  (8)
```

1. Loads the mrgsolve package.
2. Reads-in a model file (pk1) from the available library or can be substituted by a user-defined model code. The user-defined model syntax resembles the one of NONMEM, and as such includes, for example, the definition of model parameters (fixed effects and interindividual, interoccasion, or residual unexplained variability), ODEs—if needed—describing the kinetics in a variety of compartments, and a specification of the relevant output, including calculation of derived parameters if desired.
3. Prints overview of the model (mod).
4. Prints parameters from the model.
5. Defines interventions or events (evnt). Only dose records are defined in this first example, but other events affecting model outcome can be inserted. An externally prepared data set can also be read-in instead. In this example, 10 dose records (current dose + 9 additional doses, addl = 9) of 100 dose units (amt = 100) with an interdose period of 24 units (ii = 24) are simulated, without additional non–dose-event records.

6. Performs simulations. Model (mod) and events (evnt) are the arguments of the `mrgsim` function that performs the actual simulation. The timeframe for observation timepoints of PK or PD compartments is defined. In this example, evaluation of the output is requested every 0.1 time units (delta = 0.1) up to and including 480 time units (end = 480) after the start.
7. Prints high-level overview of results (out).
8. Plots the results. Output can also be further processed using default R coding.

A full description of the model syntax as required by mrgsolve is outside the scope of this tutorial. A description of this is available in supporting documents online (https://mrgsolve.org/vignettes/), and a more extensive example is available in a published tutorial.[21]

The stepwise build-up of the adaptive simulation framework is presented from simple to complex examples. First, the basic principles of adaptive dosing in mrgsolve are presented for a single subject where a dose is modified when the plasma concentration (CP) goes above a specific threshold[22] [Adaptive_simple]. This can be evaluated either continuously (i.e., first occurrence) or evaluated at a pre-specified timepoint (Example 1). This can be further expanded to a simulation in which the dose modification takes place within a specific timeframe (Example 2). Beyond this, we continue to show how to include various components (e.g., PK, PD, and safety models) and how a safety marker can drive a dose modification (Example 3). Subsequently, more complex frameworks are exemplified in which the dose is stopped and/or reduced based on regular evaluations of more than one safety marker (Example 4) or with evaluations at various time intervals (Example 5). In the final example, a full adaptive dose simulation framework including interindividual and interoccasion variability, covariate effects, and residual unexplained variability using a more complex framework of decision rules is laid out (Example 6).

Example 1: Continuous evaluation and evaluation at a pre-specified timepoint

In this example, the PK follows a simple one-compartment kinetics with first-order oral absorption, bioavailability (fbio) fixed to 1, and a first-order elimination rate. In this basic example, a typical subject is simulated and thus any form of variability is excluded. The dose is reduced by 50% for all future doses once a threshold is reached.

As shown in the ***R control script continuous and at pre-specified timepoint***, a data set [data] is needed, which contains the minimal input features that need to be provided to mrgsolve (see supplemental material for the practical implementation):

- The predetermined dose regimen [`dose`] using events [`ev()`].
- When evaluated continuously (i.e., at the first occurrence; Example 1a), nothing else needs to be defined in the data. If the evaluation is done at a pre-specified timepoint (Example 1b), the definition of the output upon which the condition is evaluated should be supplied [`evalt`]. In this basic example, the evaluation is performed at a fixed timepoint of 74 h (time = 74) for which an event identifier (EVID==33) is provided. Note that this identifier can be set to any nonreserved value but can be intentionally numbered to be linked to a certain chain of events to support bookkeeping. For observation events, the parameters amt, ii, addl, and cmt are included to correctly combine events but will be ignored in the model code. ID represents the subject identifier, and cmt is the observed compartment.

R control script continuous and at pre-specified timepoint

```
dose <- ev(amt = 1200, ii = 12, addl = 19)
evalt <- data.frame(time=74,amt=0,ii=0,addl=0,
cmt=1,evid=33,ID=1)
data <- rbind(as_data_set(dose),evalt)
print(data)

#   time  amt ii addl cmt evid ID
# 1    0 1200 12   19   1    1  1
# 2   74    0  0    0   1   33  1
```

As shown in the ***mrgsolve model coding continuous and at pre-specified timepoint***, the model specification of the PK model (one-compartment, oral administration) with the required parameters (firs-torder absorption rate constant: KA, clearance: CL, volume of distribution: V, fbio = 1; note that F_DEPOT is a reserved variable for a library PK model) has:

- The condition value at which an intervention needs to take place (CP = 100) [`condition`].
- A logical variable keeping track whether the condition is met [`condition_met`]. The condition can be either evaluated continuously or at a pre-specified timepoint (i.e., EVID==33).
- It is stated what is the type of intervention that needs to take place when the condition is met (50% reduction of the dose via `dredf`). In the code in the supplemental material, it can be seen that `dredf` is pre-specified in the $PARAM line. This allows to control its value outside the model code. Note that in subsequent examples `dredf` is specified not in the $PARAM line.

mrgsolve model coding continuous and at pre-specified timepoint

```
$PK
if(condition_met) F_DEPOT = fbio * dredf;

// continuous evaluation of the condition
$ERROR
condition_met = CP > condition || condition_met;

// evaluation of the condition at a pre-specified
timepoint
$ERROR
condition_met = (EVID==33 && CP > condition) ||
condition_met;
```

The resulting output of a continuous evaluation of the condition (Example 1a, i.e., without coding the EVID==33 in the `condition_met`) is shown in the upper panel of Figure 3. Construction of these types of plots displaying conditions over time and the intended dose reduction factor are highly recommended to check if coding is correctly implemented but also to visualize the previously chosen criteria and facilitate intermediate team discussions.

Once the CP reaches a value of 100, the condition is met (`condition_met` goes from false to true), and the dose is reduced by half for all subsequent doses. Note that if the CP were to go above 100 a second time, there would not be an additional dose reduction; in this basic example, the condition has already been met and the dose does not change to a lower value.

The lower left panel shows the case where the CP is measured only at a pre-specified date or visit (Example 1b, i.e., at EVID==33). This behavior is included by (1) adding flagged visit events in the mrgsolve command section and (2) adding the event flag to the coding of the reduction in dose (`dredf`). As a result, the dose is changed only after the CP and event condition are met. This can be observed in the plot, where the dose reduction factor is changed after the pre-specified visit, even though the threshold is reached earlier.

Example 2: Evaluation within a pre-specified timeframe

Example 1b can be expanded to a situation in which the dose reduction takes place within a specific timeframe (Example 2). In this case (as shown in the ***R control script within a pre-specified timeframe***), additional event records are included with a separate identifier to mark the end of the evaluation period (EVID==34 in `stopt`). This extra event record allows the inclusion of an extra condition in the model code, which can be used to restart the dose. Note that this event is set to 131.9 and not 132 (e.g., just before the new dose is administered). This is done for all future code examples and is an important technical requirement to track sequence of events. Here this is hardcoded for illustration purposes, and it is independent of the plasma concentration rising above the threshold value (see Figure 3, lower right panel). The ***additional mrgsolve model coding within a pre-specified timeframe*** is shown below.

R control script within a pre-specified timeframe

```
dose <- ev(amt = 1200, ii = 12, addl = 19)
evalt <- data.frame(time=74,amt=0,ii=0,addl=0,
cmt=1,evid=33,ID=1)
# note that time of reset is just before upcoming
dose
stopt <- data.frame(time=131.9,amt=0,ii=0,ad-
dl=0,cmt=1,evid=34,ID=1)
data <- rbind(as_data_set(dose),evalt,stopt)
print(data)

#   time  amt  ii   addl cmt evid ID
# 1   0.0 1200 12   19   1    1   1
# 2  74.0    0  0    0   1   33   1
# 3 131.9    0  0    0   1   34   1
```

Additional mrgsolve model coding within a pre-specified timeframe

```
$PK
if(condition_met) F_DEPOT = fbio * dredf;

$ERROR
condition_met = (EVID==33 && CP > condition) ||
condition_met;
if(EVID==34) condition_met=false;
```

Example 3: Inclusion of multiple effect markers and evaluation based on one

Example 3 aims to show how to include various components (e.g., PK, PD, and safety) and how a safety marker can drive a dose modification. In oncology, this applies to, for example, safety events, or combinations thereof, that lead to dose modifications, for instance, when myelosuppression occurs, or when fatigue or vomiting events happen. From the simulations, the frequency of dose changes, the steady-state dose, the predicted PD profile, and any other predictions can be derived to help understanding the impact on overall treatment usage if dose changes are triggered by one or more (safety) end

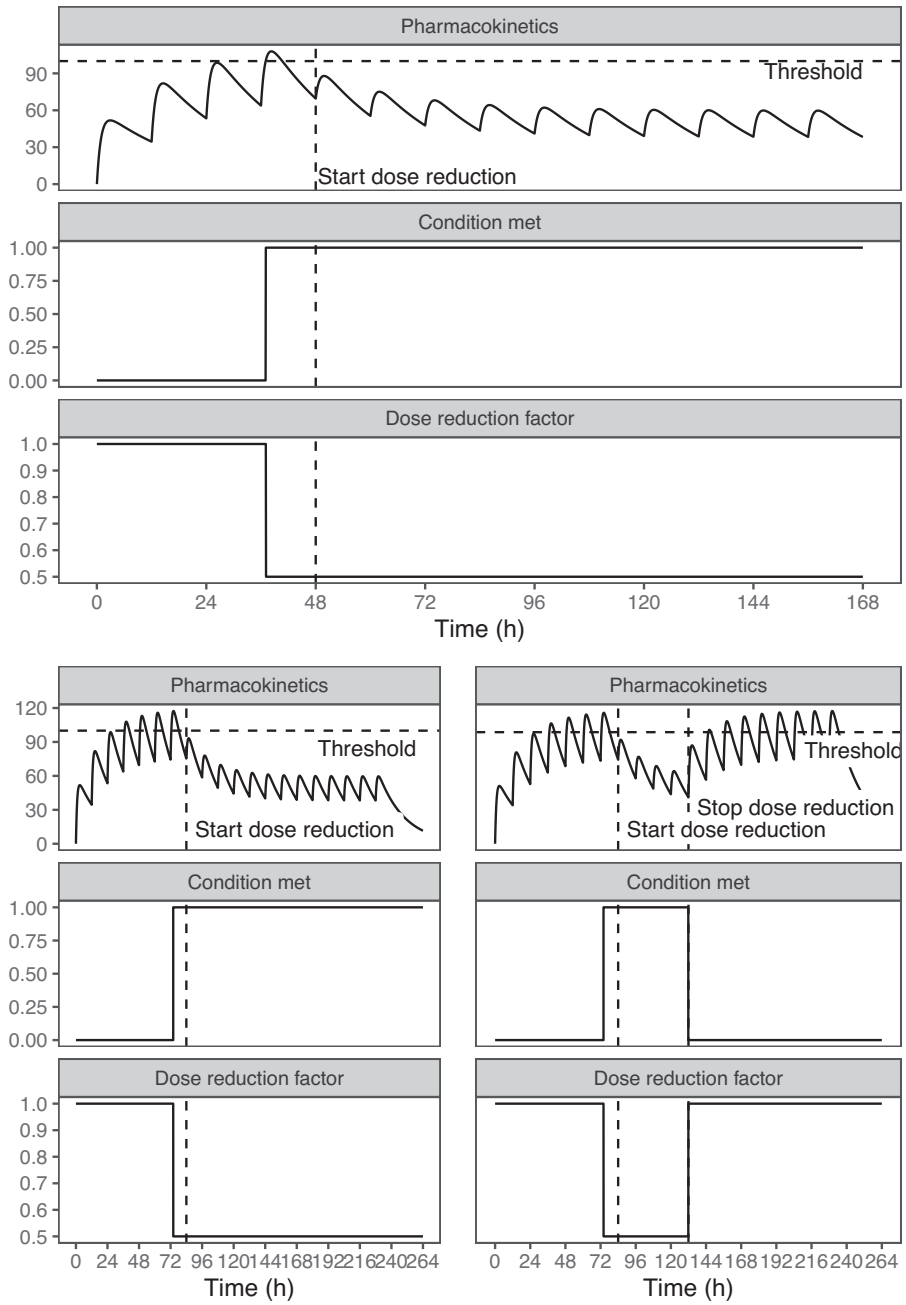

FIGURE 3 Simulated typical concentration-time profile over time for the basic Examples 1a, 1b, and 2. A dose reduction occurs if the condition of CP = 100 is met (1) at the first occurrence (upper panel) or (2) evaluated at a specific visit only (lower left panel) or (3) evaluated within a pre-specified timeframe (lower right panel). Note that there is no variability included in these basic examples.

points. Visualization of the combined results also becomes important to show this impact on overall treatment usage.[23]

In the basic examples, the dose adaptation element was solely based on an exposure measure and/or a specific timing or timeframe. However, in practice it might not just be PK thresholds or metrics that can prompt a dose adaptation, but combinations of PD and/or safety markers as well. If there are suitable models to describe the PK, PD, and safety markers in the intended population, these can be used to define dose adaptation conditions relative to certain thresholds or events. In such a setup, it should be noted that all separate models are preferably combined into a single mrgsolve model. This ensures that changes in one marker are directly reflected in all corresponding markers during the simulation.

Example 3 presents a single subject where the framework for dose adaptations includes a combination of PK, a PD marker (target engagement marker), and a safety marker (e.g., myelosuppression). Again, a typical

subject is simulated, and any form of variability is not included.

In this example, a "safety marker 1" or SM1 is used as a criterion for a temporary dose interruption. Critical aspects of this design are as follows:

- Final model parameter estimates are set within the R control script [`parm in` ***R control script multiple markers and evaluations***].
- Event records are made for dosing [`evnt`], SM1 evaluation after 3 weeks [`evalt`], and possible dose restart after 4 weeks [`stopt`].
- A safety threshold minimum for SM1 [`sm1thresh` in the ***Additional mrgsolve model coding multiple markers and evaluations***] is defined based on clinical standards and added to the mrgsolve model. When the SM1 reaches a value below this level at a pre-specified timepoint (at EVID==33), the drug dosing is interrupted.
- The dosing is temporarily interrupted via `dref`.
- After an interruption of 1 week (168 h), dosing is automatically resumed (at EVID==34).

R control script multiple markers and evaluations

```
parm <- c(CL = 60, VC = 250, KA = 2.5, Q = 11,
VP = 440,
        TEM_BSL = 100, KOUT = 8e-3, EC50 =1.5, EMAX
        = -0.8,
        MTT_SM1 = 130, CIRC0_SM1 = 240, GAMA_SM1 =
        0.3, SLOPE_SM1 = 4e-4)

evnt <- as_data_set(ev(amt = 400*1000, ii = 24,
addl = 50,tinf=3.5))
evalt <- data.frame(time=504-0.01,amt=0,rate=0,
ii=0,addl=0,cmt=1,evid=33,tinf=0,ID=1)
stopt <- data.frame(time=672-0.01,amt=0,rate=0,
ii=0,addl=0,cmt=1,evid=34,tinf=0,ID=1)
evnt2 <- rbind(evnt,evalt,stopt)
evnt2

#     time    amt      rate ii addl cmt evid tinf ID
# 1   0.00 4e+05 112359.6 24   50   1    1  3.5  1
# 2 503.99 0e+00       0.0  0    0   1   33  0.0  1
# 3 671.99 0e+00       0.0  0    0   1   34  0.0  1
```

Additional mrgsolve model coding multiple markers and evaluations

```
$PK
F1 = fbio * dredf;

$ERROR
if(EVID==33 && A(9) < sm1thresh) dredf = 0;
if(EVID==34) dredf=1;
```

The resulting behavior is illustrated in Figure 4. The upper panel displays a subject without events; the SM1 levels remain above the defined threshold and dosing is maintained, whereas in the lower panel, the value for SM1 at the pre-specified timepoint is below the threshold and dosing is thus withheld for a fixed period of 1 week. SM1 values rise after dosing is withheld, but even if they do not reach the threshold, dosing is restarted anyway because the return to normal dosing is hardcoded to occur after one week of interruption. For both instances, the resulting profiles of the PD marker (target engagement marker) are also presented in Figure 4.

Example 4: dose stop and reduction with multiple evaluations based on more than one safety marker

A positive aspect of this type of framework is its adaptability. It allows testing for more complex dose modifications driven by one or more (safety) markers. Other more complex options for dose reductions or dose interruptions are possible, such as dose hold until one or more markers are back above their threshold, or extent of dose reduction based on prior dose level and marker values at time of dose hold. This is useful, for instance, to get insight into the impact on overall treatment usage if dose changes are triggered to establish dosing recommendations in relation to safety events. An example of such an intricate dose modification schedule based on a combination of dose and platelet count is included in the prescribing information for ruxolitinib, which could have been evaluated during development with such a framework.[24]

In Example 4, a more elaborate implementation is presented with the combination of a 1-week dose stop due to either SM1 and/or safety marker 2 (SM2) dropping below their predefined thresholds together with a more complex restarting protocol. During the stop week, a reassessment is done every 48 h for both safety markers to evaluate whether values return above the threshold. The status of both safety markers is also evaluated after a week. If after a week one of the safety markers is below its threshold, the dose is withheld until these criteria are met. If both the SM1 and SM2 values returned to safe levels after a week, but at two subsequent 48-h evaluations during the week the criteria were not met, dosing is restarted with a 50% reduction in dose. In case the initial dose could be restarted, independent of the two markers returning to safe levels, an additional event ID should be implemented for dose restarts. For the current example, this is not done to keep it simple. However, this will be shown in subsequent examples.

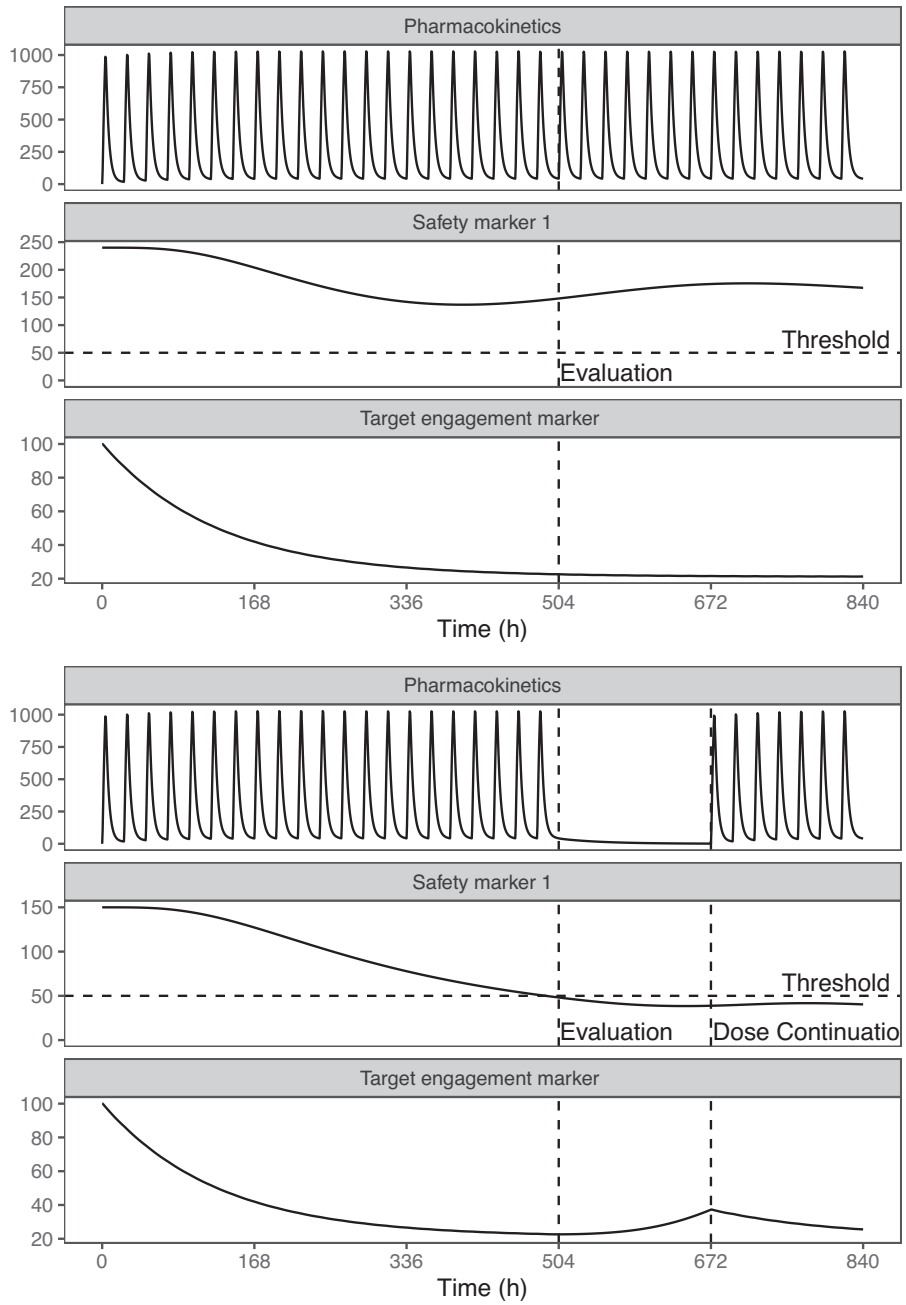

FIGURE 4 Simulated typical pharmacokinetics, safety, and target engagement profiles over time for the framework showing a subject without dose adaptation (upper panel) and a subject with a temporary 1-week dose interruption prompted by Safety Marker 1 (lower panel). Note that there is no variability included in this example (Example 3).

Critical aspects of this design are as follows:

- Separate event records are made in the ***R control script dose stop and reduction*** for dosing [`evnt`], for weekly dose stop/restart events [`dose_stop`], and events are made for the three evaluation moments (48, 96, and 144 h) after the start of the dose interruption week [`eval_a, eval_b, eval_c`].
- Safety thresholds are set for the SM1 [`sm1thresh`] and for SM2 [`sm2hresh`]. In the ***mrgsolve model coding dose stop and reduction***, when one or both safety makers are below their respective threshold at a pre-specified timepoint (at EVID==33), the drug dosing is interrupted [`dose_stop`], and the dosing is temporarily interrupted via setting `dredf` to zero.
- During the 1-week interruption [`dose_stop==true`], the levels of the safety markers are evaluated every 48 h (at EVID==34, EVID==35, and EVID==36). Depending on whether one or both safety markers are

(still) below their respective threshold, the evaluation events are set to true [eval_a, eval_b, eval_c].
- After week (at EVID==33), the status of the safety markers is read out. If two subsequent evaluation events [eval_a **and** eval_b] or [eval_b **and** eval_c] had values below their set thresholds, dosing is continued (given the criteria discussed previously) after a week with a reduced dose [dose_red==true; if(dose_red) dredf=0.5]. If after a week both safety makers are above their thresholds, the previous dose level is continued [dredf ==1].

Note that the EVID flags are included at set times, but whether these flags are used is dynamic depending on the outcome of safety events. Nevertheless, these flags are required for bookkeeping purposes. Note that this example includes sampling of the residual error for both the SM1 and SM2 levels to mimic actual measurements (see Appendix S1). During the evaluation moments, the simulated safety maker values include the residual error [SM1_RE, SM2_RE] and are compared against their respective thresholds. Interindividual and interoccasion variability is sampled either in R or in the mrgsolve code to enable simulation of different individuals.

R control script dose stop and reduction

```
evnt <- as_data_set(ev(amt = 400*1000, ii = 24,
addl = 50,tinf=3.56,realize_addl=TRUE))
dose_stop    <- data.frame(time=
seq((7*24),(50*24),7*24)-0.001,
                           amt=0,rate=0,ii=0,
addl=0,cmt=1,evid=33,tinf=0,ID=1)
eval_a       <- data.frame(time=
seq((7*24)+48,(50*24)+48,7*24)-0.001,
                           amt=0,rate=0,ii=0,
addl=0,cmt=1,evid=34,tinf=0,ID=1)
eval_b       <- data.frame(time=
seq((7*24)+96,(50*24)+96,7*24)-0.001,
                           amt=0,rate=0,ii=0,
addl=0,cmt=1,evid=35,tinf=0,ID=1)
eval_c       <- data.frame(time=
seq((7*24)+144,(50*24)+144,7*24)-0.001,
                           amt=0,rate=0,ii=0,
addl=0,cmt=1,evid=36,tinf=0,ID=1)
evnt         <- rbind(evnt,dose_stop,eval_a,eval_b,
eval_c)
```

Additional mrgsolve model coding dose stop and reduction

```
$PK
F1 = fbio * dredf;

$ERROR
// First perform the weekly evaluations and set
the dose_stop flag if necessary (values are below
threshold)
dose_stop = ((EVID==33) && (SM1_RE < sm1thresh |
SM2_RE < sm2thresh)) || dose_stop;

// Perform the evaluations if dose stops are in
place
if(dose_stop==true){
  eval_a     = ((EVID==34) && (SM1_RE < sm1thresh
| SM2_RE < sm2thresh)) || eval_a; // first evalua-
tion moment
  eval_b     = ((EVID==35) && (SM1_RE < sm1thresh
| SM2_RE < sm2thresh)) || eval_b; // second evalu-
ation moment
  eval_c     = ((EVID==36) && (SM1_RE < sm1thresh
| SM2_RE < sm2thresh)) || eval_c; // third evalua-
tion moment
}

// Reduce or stop the dose based on the outcome
if(dose_stop==true) dredf = 0;
if(EVID==33 && ((eval_a==true & eval_b==true) |
(eval_b==true & eval_c==true))) dose_red = true;
// Notice here that dose is only restarted in case
values are above threshold, otherwise dose is
withheld
if(EVID==33 && (SM1_RE > sm1thresh & SM2_RE > sm-
2thresh) && dredf==0){
  dose_stop = false;
  dredf = 1;
  if(dose_red) dredf = 0.5;
   }
```

A potential outcome of this example is illustrated in Figure 5. The output for SM1 is such that its value drops below the threshold after 2 weeks, and consequently dosing is withheld. During the week of the dose stop, the three evaluations show that the safety markers do not rise above their threshold at two subsequent 48-h evaluations. At the weekly evaluation they are above their threshold, so the dose can be restarted, but the outcome of the 48-h evaluations results in a dose reduction as seen in the dose reduction factor.

Example 5: weekly versus three weekly evaluations of end points

In the previous example, evaluations were done every week. Another option is to initially have an evaluation every week for the first 3 weeks and subsequently shift the evaluation to once every 3 weeks. Within this design (Example 5), a

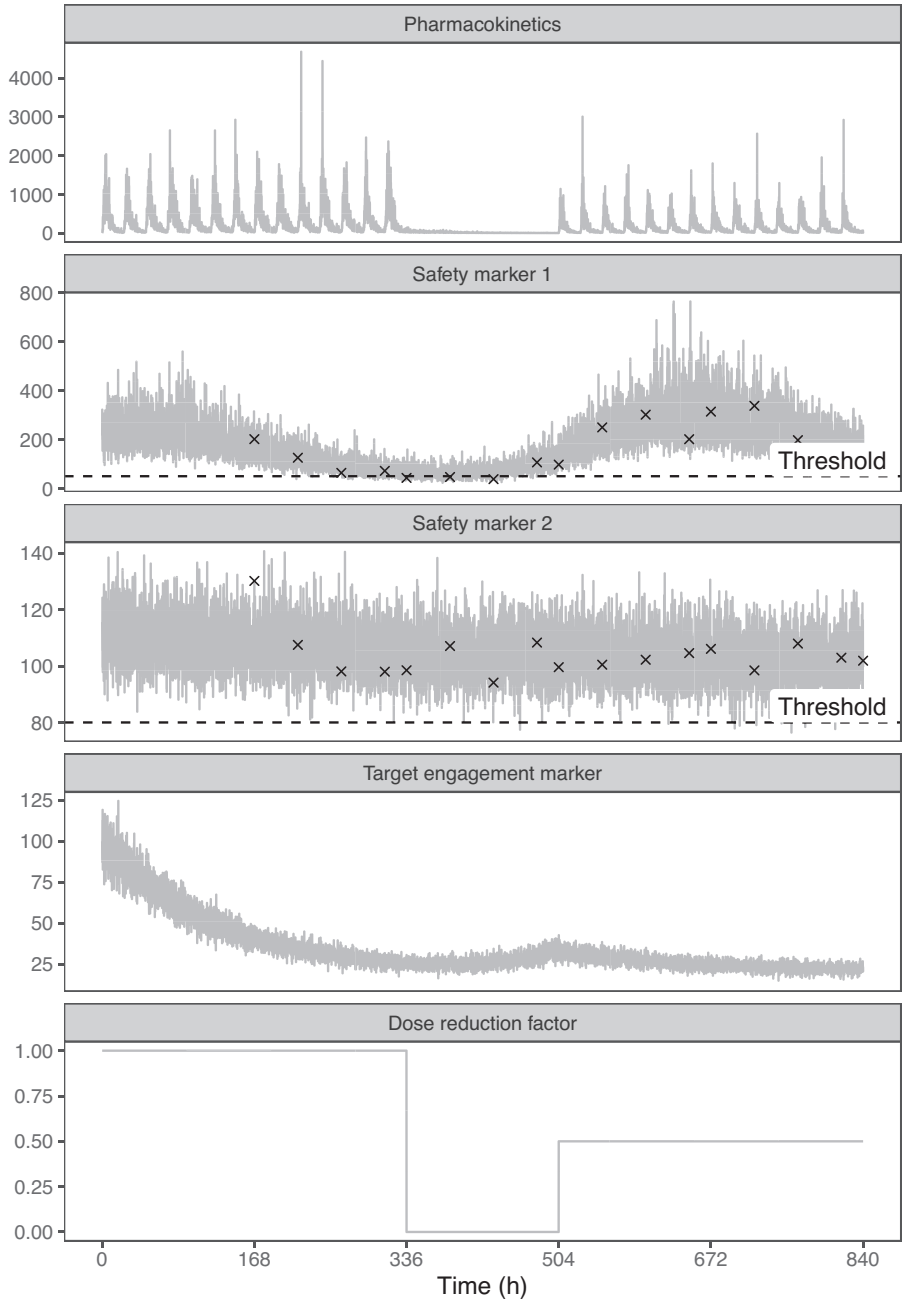

FIGURE 5 Simulated pharmacokinetics, safety, and target engagement profiles over for the advanced framework, with scheduled evaluations within a fixed dose interruption period based on more than one safety marker (Example 4). This example includes interindividual, interoccasion, and residual unexplained variability. A single subject is shown here. The dense output of the simulation is shown in gray, and the sampling at the evaluation moments for the safety markers are presented by the black symbol (×).

counter is used to enable evaluations every 3 weeks, which still requires an event every week (EVID==33) to keep track of a potentially dynamic dose interruption rhythm. Although this seems like a simple addition to Example 4, the start of the "new" period can deviate based on the occurrence of a dose stop at Week 3. To do this, multiple flags are used and evaluated. The first flag tracks whether the three weekly evaluations have started [`threewkeval` in ***Additional mrgsolve model coding weekly*** `versus` ***three weekly evaluations***]. The second flag keeps track of the week number [`weekno`], which is important to initiate the three weekly evaluations and to shift the three weekly evaluations based on any dose stop event during the weekly evaluations and/or thereafter.

Critical aspects of this design are as follows:

- The three-week switch [`threewkeval`] starts after three (`TIME==504 & ...`) or after four (`TIME==672`) weeks depending on a dose stop (... `& F1!=0`) at Week 3.

- During the first 3 weeks [`if(threewkeval==-false)`], there is a weekly evaluation (at EVID==33) to check whether the SM1 or SM2 drop below their respective thresholds. If so, the dose is stopped [`dose_stop`] via `dredf` as in Example 4 with an evaluation every 48 h. If there is a dose stop during the first or second week, this does not influence the start of the three weekly evaluations; if there is a dose stop required at the third week, this will shift the start of the three weekly evaluations by 1 week, that is, [`threewkeval`] remains false and is reevaluated at Week 4 (`TIME==672`). Either way, at Week 4 the previous dose will be restarted or a dose reduction takes place.
- Once the three-weekly evaluation starts, the week number [`weekno`] is tracked to indicate when the 3 weeks ended to check the status of the safety markers. If more than 3 weeks have passed or when a dose stop is required during this three-weekly period, this counter is reset to 1, or the former to again track 3 weeks, for the latter to be able to shift the three weekly evaluations relative to the required 1-week dose stop. Notice that the [`threewkeval`] remains true and is not reset.

R control script weekly versus three weekly evaluations

no additional changes needed to the R code compared to Example 4

Additional mrgsolve model coding weekly versus three weekly evaluations

```
$ERROR
// the three-week switch starts after three or
after 4 weeks depending on a dose stop at week 3
if((TIME==504 & F1!=0) || TIME==672)
threewkeval=true;

// handling of weekly or 3 weekly evaluations is
split out
if(threewkeval==false){
  dose_stop = ((EVID==33) && (SM1_RE < sm1thresh |
SM2_RE < sm2thresh)) || dose_stop;
}else{
  if(EVID==33) weekno = weekno &#x002B; 1;
  // this is the actual time for evaluations/dose
stops - take into account the sequence here!!
  if(EVID==33 & weekno > 3) dose_stop = ((EVID==33)
&& (SM1_RE < sm1thresh | SM2_RE < sm2thresh)) ||
dose_stop;
  if(weekno>3) weekno = 1;
  if(dose_stop) weekno = 1;
}
```

Figure 6 presents a potential outcome for this example. During the first 3 weeks, the safety markers are evaluated each week. At Week 3 (TIME==504), the SM1 drops below its threshold and the dose is stopped. As a result of the dose stop, the three weekly evaluations do not yet start as F1 is effectively set to zero. A week thereafter, at Week 4 (TIME==672), the three weekly evaluations do start, and the safety markers are evaluated again. At Week 5 (TIME==840), the safety markers are again evaluated and are both above their thresholds with two subsequent 48-h evaluations also above the thresholds. Therefore, the previous dose is given. Notice that subsequently the week-number counter is increasing every week. Two weeks later, the dose is stopped as SM1 drops below its threshold and based on the outcome of the 48-h evaluations, at Week 8 (TIME==1344) there is again an evaluation moment, which now results in a dose reduction, that is, at the weekly evaluation the levels are above their threshold, but two subsequent 48-h evaluations were not.

Example 6: a full adaptive dose simulation framework including variability, covariates, and residual unexplained variability

A possible full adaptive dose simulation framework including variability, uncertainty, covariates, and residual unexplained variability is presented in Example 6. It combines the PK, PD (target engagement marker), and three safety markers. Fukae et al present an example, including graphs summarizing the outcome, where the exposure–response for multiple efficacy and safety end points was used to justify the clinical dose of valemetostat for adult T-cell leukemia/lymphoma.[25] The models for PK, PD, and safety can contain expressions for within- and between-subject variability, covariate relationships, parameter uncertainty, and residual unexplained variability. In the supplemental material, mrgsolve and R codes are provided. This framework combines multiple technical elements from the previous examples with the set of decision rules visualized in Figure 7.

These decision rules in Figure 7 illustrate that within the first 3 weeks, SM1 and SM2 are evaluated weekly to assess whether they are below their respective threshold values (Evaluate). After 3 weeks, the frequency of these evaluations is changed to once every 3 weeks and an assessment for Safety Marker 3 (SM3) is done, for which events are described using a logistic function. If somewhere during the first 3 weeks a safety event occurred for this third marker, the dose is stopped for a fixed duration of 1 week (one-time fixed 7-day interruption). If SM1 and/or SM2 drop below their respective threshold during the weekly evaluations, the following is done: the dose is

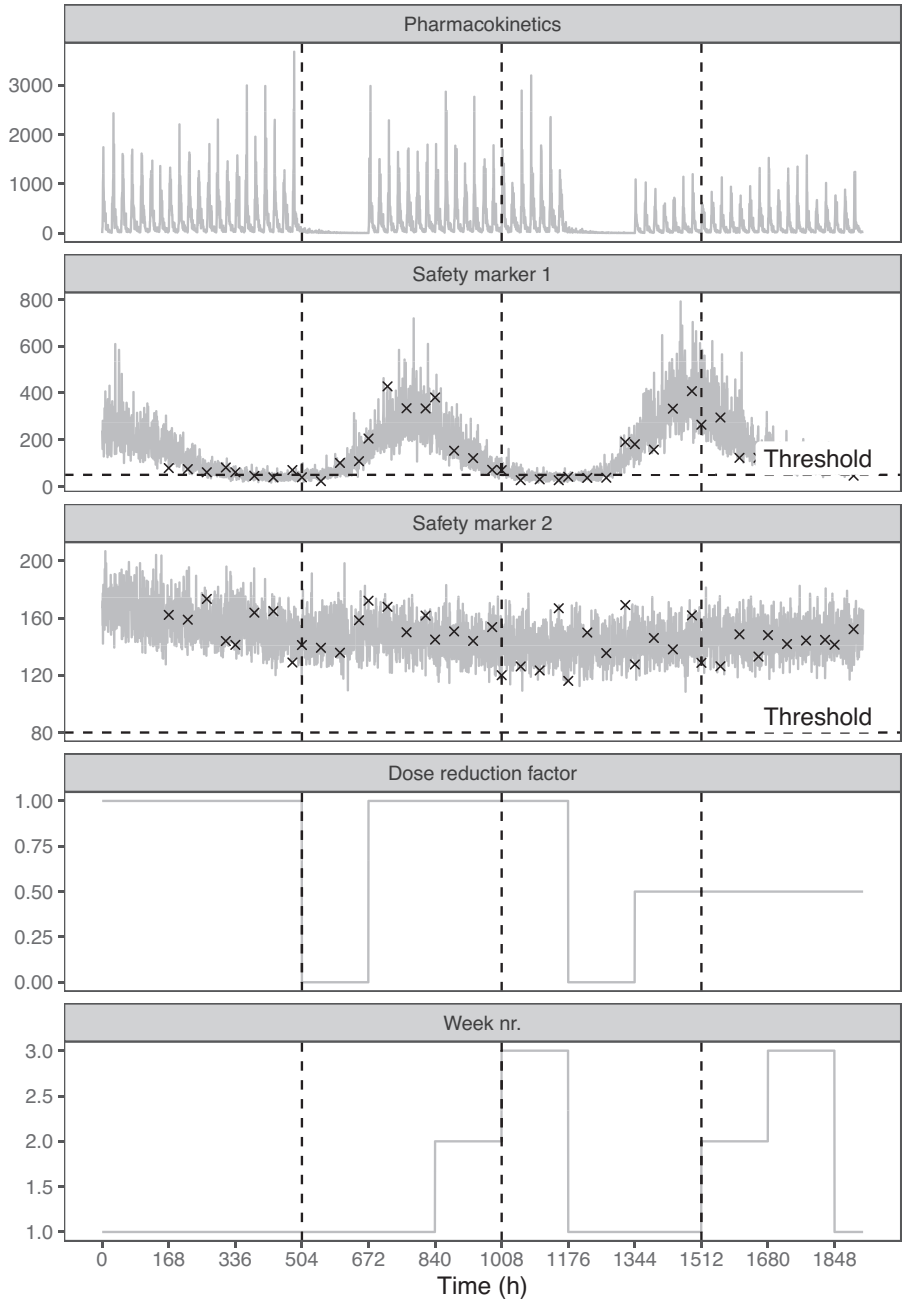

FIGURE 6 Simulated pharmacokinetics, safety, and target engagement profiles over time for the advanced framework, with weekly and three weekly evaluations depending on the status of the safety markers (Example 5). This example includes interindividual, interoccasion, and residual unexplained variability. A single subject is shown here. The dense output of the simulation is shown in gray, and the sampling at the evaluation moments for the safety markers are presented by the black symbol (×); Week nr, week number.

stopped, and during that Week 3 evaluation moments are performed every 48 h (fixed 7-day interruption[1]). If two consecutive evaluations during this week remain below threshold, the dose is stopped for another week (fixed 7-day interruption[2]), otherwise dosing is continued with the same dose. When the dose is stopped for another week, three evaluations are again done every 48 h. If at the end of the second dose stop Week 2 consecutive measurements during that week remained below threshold, the treatment is stopped (stop treatment). The decision rules in Figure 7 show that if two consecutive measurements are above threshold after that second week, the dose is reduced and restarted (first or second dose reduction). This dose reduction can take place twice.

Various complexities are implemented here that can be found in the supplemented code. The example is not discussed in further detail, but the important learning is to have events within the event data set to account for all possible options that may occur. Herein, the sequence of events is extremely important. Furthermore, the events

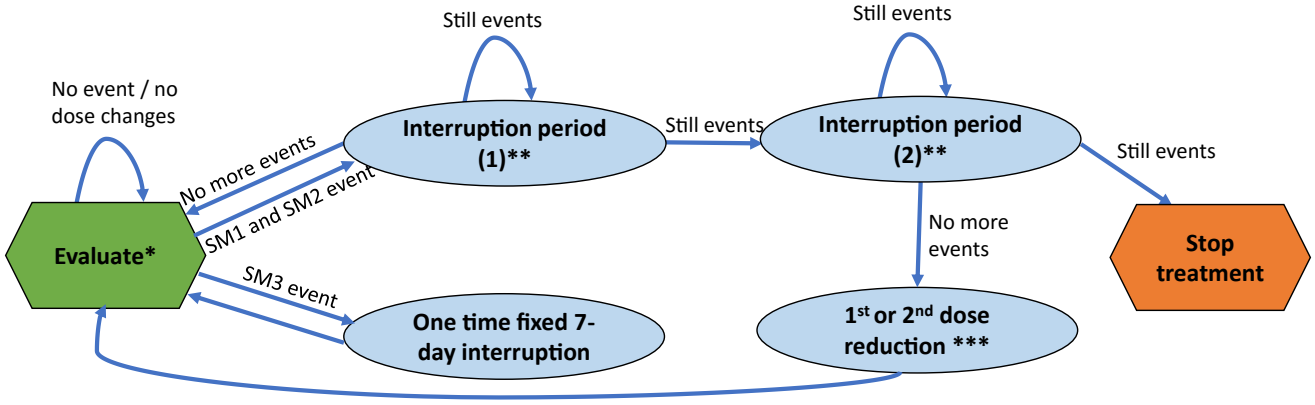

FIGURE 7 Decision rules for a full adaptive dose simulation framework. *Evaluations for Safety Marker 1 (SM1) and Safety Marker 2 (SM2) are done weekly during the first 3 weeks, after that evaluations are done once every 3 weeks. Safety Marker 3 (SM3) is evaluated once after 3 weeks. **Within the interruption period (1 week), a sample is taken at 48 h. If event is resolved, treatment continues; if not, a second and/or a third sample is taken. If this continues for more than a week, the dose is reduced or treatment is stopped. ***First dose reduction, and after dose has been reduced once there can be a second reduction.

should be picked up by the model to implement the logic described previously. Within the supplied R script there are plots coded to look at the dynamics of the various markers and event switches. In the code provided in the supplemental material, the sequence with respect to timing is important, for example, does it occur just before or after a dose and if two events take place at the same time, does one have to be set before the other? Note that the numbering in general is arbitrary and only important for bookkeeping, although including a logical numbering greatly aids understanding the sequence of events. The following events are defined:

- 33: weekly evaluation for dose stop (just before dosing)
- 34: reset for weekly evaluation for dose stop (just before dose stop evaluation)
- 35: Evaluation Moment 1 during dose stop (just before dosing)
- 36: Evaluation Moment 2 during dose stop (just before dosing)
- 37: Evaluation Moment 3 during dose stop (just before dosing)
- 38: reset for previous dose stop and evaluation moments (just after dosing)
- 39: evaluation of first dose reduction (before dosing and weekly evaluations)
- 40: evaluation of second dose reduction (before dosing and weekly evaluations)
- 41: evaluation of dropout due to SM3

The code could be further expanded to have an evaluation of the dose stop every 48 h instead of weekly and, depending on the outcome, the dose could then be reduced or continued immediately. This would present a more dynamic design in which treatment is not withheld for a week per se.

CONCLUDING REMARKS

The adaptive dose simulation framework offers the opportunity to explore the impact of starting dose and event-driven dose and regimen changes on clinical end points. One of the advantages of such a framework lies in that it supports timely and informed decisions (such as for dose selection) by using early clinical data in an integrated assessment. Moreover, it allows the exploration of dosing regimens (for both monotherapy and combinations) that potentially better balance benefit and (chronic) adverse effects and find effective dose modification strategies that minimize the unscheduled dose interruptions, and it can provide insights on whether special populations may need different dosing strategies. As such, in silico exploration can prioritize dosing regimens for future testing and thus reduce the risk for patients to be exposed to regimens that either provide limited benefit or unacceptable risks. Moreover, it can provide insights into the steady-state dose requirements and the utilization of the various dose strengths that can be useful for clinical supply and manufacturing planning. This adaptive dose simulation framework does have a broad applicability in the context of dose optimization, in particular for, but not limited to, compounds that have a narrow therapeutic index, such as often is the case in the oncology space.[6,26] With the highly needed reformation of the dose optimization and dose selection paradigm in oncology drug development, such a framework can help focus on maximizing clinical benefit and minimizing (chronic) adverse effects.[3–7]

One approach to maximize the opportunity for therapeutic benefit is to ensure that patients do have sufficient exposure to the drug by measuring the drug exposure through therapeutic drug monitoring. This is particularly relevant for compounds that may have large PK variability

or a narrow therapeutic index. Well-known examples are digoxin in heart failure, cyclosporin after organ transplantation, and phenytoin for seizures. Although not commonly used, it has also been investigated in oncology for kinase inhibitors.[27,28] In other cases, the dose may need to be carefully titrated in patients to carefully attain the desired effect and/or to minimize adverse effects. An example is the individualized hemoglobin-level guided dosing strategy to treat anemia with erythropoietin stimulating agents or Hypoxia-inducible factor inhibitors in which the hemoglobin baseline and the rate of change during treatment determine the starting dose and dose modifications, respectively.[29]

This tutorial has introduced various general and technical aspects that could be part of an adaptive dose simulation framework, which could aid development of compounds benefiting from dose modifications to manage tolerability or address suboptimal efficacy. The examples should be considered as a starting point and are appropriate to investigate drugs based on (longitudinal) changes in their efficacy and safety profiles and can prompt dose switches or (temporary or permanent) dose interruptions. Basic to more advanced examples are presented, highlighting the flexibility to combine multiple end points (PK/PD/safety), which can lead to dose adaptations. In addition, models for PD or safety end points that are included in the simulation framework may need to be based on an appropriate analysis that takes into account actual dose modifications during treatment.[30]

Beyond the impact that the dose simulation framework has on the selection of dosing regimens and exploring effective dose modification strategies, the practical benefit is that it enhances team engagement and optimizes the extraction, and combination, of information from various sources. With that, professionals with diverse backgrounds and interests are brought to a similar (minimum) knowledge level and discuss outcomes in an integrated fashion. There also lies the opportunity in the ability to (re-)use models and/or to inform these models based on (pre)clinical data ahead of clinical trials. Information is captured and attained knowledge is preserved. Considerable time is saved for the user, by implementation of standard scripts and code, which facilitates timely support to the decision-making. Furthermore, the ability to provide insight of the influence of (clinical) decision criteria (vs.) compound characteristics/pharmacology (and vs. other drugs and/or combinations of drugs) will provide earlier insight into what development directions might be beneficial. Ultimately, such a framework could assist in benefit to the patient, such as early dose optimization (also in combination setting) or individualization of dosing requirements for special populations.

ACKNOWLEDGMENTS

We thank Maurice Ahsman for setting up an initial draft of this tutorial.

FUNDING INFORMATION

Funding was provided by GSK.

CONFLICT OF INTEREST STATEMENT

R.H., E.M., and T.M.P. are employees of Leiden Experts on Advanced Pharmacokinetics and Pharmacodynamics and collaborated under a consultancy agreement with GSK on the simulation framework as described in this tutorial. R.P., G.F.-B., and S.A.G.V. are employees of and stockholders in GSK. The authors declared no competing interests for this work.

ORCID

Geraldine Ferron-Brady https://orcid.org/0000-0001-5703-8315
Teun M. Post https://orcid.org/0000-0001-9183-2678
Sandra A. G. Visser https://orcid.org/0009-0008-2186-7195

REFERENCES

1. Witt DM, Clark NP, Kaatz S, Schnurr T, Ansell JE. Guidance for the practical management of warfarin therapy in the treatment of venous thromboembolism. *J Thromb Thrombolysis*. 2016;41(1):187-205. doi:10.1007/s11239-015-1319-y
2. Roda D, Jimenez B, Banerji U. Are doses and schedules of small-molecule targeted anticancer drugs recommended by phase I studies realistic? *Clin Cancer Res*. 2016;22(9):2127-2132. doi:10.1158/1078-0432.CCR-15-1855
3. FDA website Project Optimus. Accessed May 22, 2023 https://www.fda.gov/about-fda/oncology-center-excellence/project-optimus
4. Bullock JM, Rahman A, Liu Q. Lessons learned: dose selection of small molecule-targeted oncology drugs. *Clin Cancer Res*. 2016;22(11):2630-2638. doi:10.1158/1078-0432.CCR-15-2646
5. Shah M, Rahman A, Theoret MR, Pazdur R. The drug-dosing conundrum in oncology - when less is more. *N Engl J Med*. 2021;385(16):1445-1447. doi:10.1056/NEJMp2109826
6. Fourie Zirkelbach J, Shah M, Vallejo J, et al. Improving dose-optimization processes used in oncology drug development to minimize toxicity and maximize benefit to patients. *J Clin Oncol*. 2022;40(30):3489-3500. doi:10.1200/JCO.22.00371
7. Murphy R, Halford S, Symeonides SN. Project Optimus, an FDA initiative: considerations for cancer drug development internationally, from an academic perspective. *Front Oncol*. 2023;13:1144056. Published 2023 Mar 3. doi:10.3389/fonc.2023.1144056
8. FDA draft guidance: optimizing the dosage of human prescription drugs and biological products for the treatment of oncologic disease. Accessed May 22, 2023. https://www.fda.gov/regulatory-information/search-fda-guidance-documents/optimizing-dosage-human-prescription-drugs-and-biological-products-treatment-oncologic-diseases

9. Guo Y, Haddish-Berhane N, Xie H, Ouellet D. Optimization of clinical dosing schedule to manage neutropenia: learnings from semi-mechanistic modeling simulation approach. *J Pharmacokinet Pharmacodyn*. 2020;47(1):47-58. doi:10.1007/s10928-019-09667-y
10. Bauer S, Demetri GD, Halilovic E, et al. Pharmacokinetic-pharmacodynamic guided optimisation of dose and schedule of CGM097, an HDM2 inhibitor, in preclinical and clinical studies. *Br J Cancer*. 2021;125(5):687-698. doi:10.1038/s41416-021-01444-4
11. Wheeler GM, Mander AP, Bedding A, et al. How to design a dose-finding study using the continual reassessment method. *BMC Med Res Methodol*. 2019;19(1):18. Published 2019 Jan 18. doi:10.1186/s12874-018-0638-z
12. Vandemeulebroecke M, Baillie M, Margolskee A, Magnusson B. Effective visual communication for the quantitative scientist. *CPT Pharmacometrics Syst Pharmacol*. 2019;8(10):705-719. doi:10.1002/psp4.12455
13. Fostvedt LK, Nickens DJ, Tan W, Parivar K. Tumor growth inhibition modeling to support the starting dose for dacomitinib. *CPT Pharmacometrics Syst Pharmacol*. 2022;11(9):1256-1267. doi:10.1002/psp4.12841
14. Walz AC, Van De Vyver AJ, Yu L, Birtwistle MR, Krogan NJ, Bouhaddou M. Leveraging modeling and simulation to optimize the therapeutic window for epigenetic modifier drugs. *Pharmacol Ther*. 2022;235:108162. doi:10.1016/j.pharmthera.2022.108162
15. R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing; 2021 https://www.R-project.org/
16. Baron K. Mrgsolve: simulate from ODE-based models. 2023. R Package Version 1.0.9. https://cran.r-project.org/web//packages/mrgsolve/mrgsolve.pdf
17. Beal SL, Sheiner LB, Boeckmann AJ & Bauer RJ, eds. *NONMEM 7.4 users guides*. (1989–2018). Accessed August 23, 2023. https://nonmem.iconplc.com/nonmem743/guides
18. Fidler M, Wilkins JJ, Hooijmaijers R, et al. Nonlinear mixed-effects model development and simulation using nlmixr and related R open-source packages. *CPT Pharmacometrics Syst Pharmacol*. 2019;8(9):621-633. doi:10.1002/psp4.12445
19. Wang W, Hallow KM, James DA. A tutorial on RxODE: simulating differential equation Pharmacometric models in R. *CPT Pharmacometrics Syst Pharmacol*. 2016;5(1):3-10. doi:10.1002/psp4.12052
20. Bachman F, Koch G, Bauer RJ, et al. Computing optimal drug dosing with OptiDose: implementation in NONMEM. *J Pharmacokinet Pharmacodyn*. 2023;50(3):173-188. doi:10.1007/s10928-022-09840-w
21. Elmokadem A, Riggs MM, Baron KT. Quantitative systems pharmacology and physiologically-based pharmacokinetic modeling with mrgsolve: a hands-on tutorial. *CPT Pharmacometrics Syst Pharmacol*. 2019;8(12):883-893. doi:10.1002/psp4.12467
22. Adaptive_simple. Accessed May 22, 2023. https://github.com/mrgsolve/gallery/blob/master/adaptive/adaptive_simple.md
23. Guo C, Liao KH, Li M, Wang IM, Shaik N, Yin D. PK/PD model-informed dose selection for oncology phase I expansion: case study based on PF-06939999, a PRMT5 inhibitor. *CPT Pharmacometrics Syst Pharmacol*. 2022. [published online ahead of print, 2022 Nov 16]. doi:10.1002/psp4.12882
24. Ruxolitinib label. Accessed May 22, 2023. https://hcp.jakafi.com/pdf/prescribing-information.pdf
25. Fukae M, Baron K, Rogers J, et al. Landmark and longitudinal exposure-response analyses for multiple efficacy and safety endpoints to justify the clinical dose of valemetostat for adult T-cell leukemia/lymphoma. https://www.go-acop.org/default.asp?abstract=493
26. Soltantabar P, Lon HK, Parivar K, Wang DD, Elmeliegy M. Optimizing benefit/risk in oncology: review of post-marketing dose optimization and reflections on the road ahead. *Crit Rev Oncol Hematol*. 2023;182:103913. doi:10.1016/j.critrevonc.2023.103913
27. Verheijen RB, Yu H, Schellens JHM, Beijnen JH, Steeghs N, Huitema ADR. Practical recommendations for therapeutic drug monitoring of kinase inhibitors in oncology. *Clin Pharmacol Ther*. 2017;102(5):765-776. doi:10.1002/cpt.787
28. Goutelle S, Guidi M, Gotta V, Csajka C, Buclin T, Widmer N. From personalized to precision medicine in oncology: a model-based dosing approach to optimize achievement of imatinib target exposure. *Pharmaceutics*. 2023;15(4):1081. Published 2023 Mar 28. doi:10.3390/pharmaceutics15041081
29. Vd Berg 2022 - Innovative adaptive dose simulations in NONMEM to allow simulation-based model evaluation of titration-based dosing, applied to a population dose-hemoglobin model for daprodustat. https://www.go-acop.org/default.asp?abstract=440
30. Kristensen NR, Agersø H. Dose/exposure-response modeling in dose titration trials: overcoming the titration paradox. *CPT Pharmacometrics Syst Pharmacol*. 2022;11(12):1592-1603. doi:10.1002/psp4.12863

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

How to cite this article: Hooijmaijers R, Parasrampuria R, Marostica E, Ferron-Brady G, Post TM, Visser SAG. Building an adaptive dose simulation framework to aid dose and schedule selection. *CPT Pharmacometrics Syst Pharmacol*. 2023;12:1602-1618. doi:10.1002/psp4.13027